

AD-A190 897

FAST ALGORITHMS FOR EULER AND NAVIER-STOKES SIMULATIONS
(U) INSTITUTE FOR SCIENTIFIC COMPUTING FORT COLLINS CO
B M JOHNSON 30 NOV 87 AFOSR-TR-87-1983 AFOSR-87-0133


1/1

UNCLASSIFIED

F/G 20/4

NL





UNCLASSIFIED

SECURITY CLASSIFICATION OF

AD-A190 897

DTIC FILE COPY

(2)

1a. REPORT SECURITY CLASSIFICATION DTIC ELECTE S D		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY JAN 25 1988		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE		4. PERFORMING ORGANIZATION REPORT NUMBER(S) AD	
5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR-87-1983		6a. NAME OF PERFORMING ORGANIZATION Inst. for Scientific Computing	
6b. OFFICE SYMBOL (if applicable)		7a. NAME OF MONITORING ORGANIZATION AFOSR/NM	
6c. ADDRESS (City, State, and ZIP Code) P.O. Box 1388 Fort Collins, Colorado 80522		7b. ADDRESS (City, State, and ZIP Code) AFOSR/NM Bldg 410 Bolling AFB DC 20332-6448	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION AFOSR		8b. OFFICE SYMBOL (if applicable) NM	
8c. ADDRESS (City, State, and ZIP Code) AFOSR/NM Bldg 410 Bolling AFB DC 20332-6448		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-87-0133	
10. SOURCE OF FUNDING NUMBERS		11. TITLE (Include Security Classification) Fast Algorithms for Euler and Navier-Stokes Simulations	
PROGRAM ELEMENT NO. 61102F		PROJECT NO. 2304	
TASK NO. A3		WORK UNIT ACCESSION NO.	
12. PERSONAL AUTHOR(S) Professor Johnson		13a. TYPE OF REPORT Final	
13b. TIME COVERED FROM 1/1/87 TO 9/30/87		14. DATE OF REPORT (Year, Month, Day) Nov. 30, 1987	
15. PAGE COUNT 12		16. SUPPLEMENTARY NOTATION	
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD		GROUP	
SUB-GROUP			
19. ABSTRACT (Continue on reverse if necessary and identify by block number) <p>→ An explicit flow solver, applicable to the hierarchy of model equations ranging from Euler to full Navier-Stokes, was combined with several techniques designed to reduce computational expense. The computational domain consisted of local grid refinements embedded in a global coarse mesh, where the locations of these refinements are defined by the physics of the flow. Flow characteristics were also used to determine which set of model equations is appropriate for solution in each region, thereby reducing not only the number of grid points at which the solution must be obtained, but also the computational effort required to get that solution. Acceleration to steady-state was achieved by applying multigrid on each of the subgrids, regardless of the particular model equations being solved. Since each of these components is explicit, advantage could readily be taken of the vector-and parallel-processing capabilities of machines such as the Cray X-MP and Cray 2.</p>			
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL Maj. John Thomas		22b. TELEPHONE (Include Area Code) (202) 767-5026	
		22c. OFFICE SYMBOL NM	

UNCLASSIFIED

INSTRUCTIONS FOR PREPARATION OF REPORT DOCUMENTATION PAGE

GENERAL INFORMATION

The accuracy and completeness of all information provided in the DD Form 1473, especially classification and distribution limitation markings, are the responsibility of the authoring or monitoring DoD activity.

Because the data input on this form will be what others will retrieve from DTIC's bibliographic data base or may determine how the document can be accessed by future users, care should be taken to have the form completed by knowledgeable personnel. For better communication and to facilitate more complete and accurate input from the originators of the form to those processing the data, space has been provided in Block 22 for the name, telephone number, and office symbol of the DoD person responsible for the input cited on the form.

All information on the DD Form 1473 should be typed.

Only information appearing on or in the report, or applying specifically to the report in hand, should be reported. If there is any doubt, the block should be left blank.

Some of the information on the forms (e.g., title, abstract) will be machine indexed. The terminology used should describe the content of the report or identify it as precisely as possible for future identification and retrieval.

NOTE: Unclassified abstracts and titles describing classified documents may appear separately from the documents in an unclassified context, e.g., in DTIC announcement bulletins and bibliographies. This must be considered in the preparation and marking of unclassified abstracts and titles.

The Defense Technical Information Center (DTIC) is ready to offer assistance to anyone who needs and requests it. Call Data Base Input Division, Autovon 284-7044 or Commercial (202) 274-7044.

SECURITY CLASSIFICATION OF THE FORM

In accordance with DoD 5200.1-R, Information Security Program Regulation, Chapter IV Section 2, paragraph 4-200, classification markings are to be stamped, printed, or written at the top and bottom of the form in capital letters that are larger than those used in the text of the document. See also DoD 5220.22-M, Industrial Security Manual for Safeguarding Classified Information, Section II, paragraph 11a(2). This form should be unclassified, if possible.

SPECIFIC BLOCKS

Block 1a. Report Security Classification: Designate the highest security classification of the report. (See DoD 5220.1-R, Chapters I, IV, VII, XI, Appendix A.)

Block 1b. Restricted Marking: Enter the restricted marking or warning notice of the report (e.g., CNWDI, RD, NATO).

Block 2a. Security Classification Authority: Enter the commonly used markings in accordance with DoD 5200.1-R, Chapter IV, Section 4, paragraph 4-400 and 4-402. Indicate classification authority.

Block 2b. Declassification / Downgrading Schedule: Indicate specific date or event for declassification or the notation, "Originating Agency Determination Required" or "OADR." Also insert (when applicable) downgrade to _____ on _____ (e.g., Downgrade to Confidential on 6 July 1983). (See also DoD 5220.22-M, Industrial Security Manual for Safeguarding Classified Information, Appendix II.)

NOTE: Entry must be made in Blocks 2a and 2b except when the original report is unclassified and has never been upgraded.

Block 3. Distribution/Availability Statement of Report: Insert the statement as it appears on the report. If a limited distribution statement is used, the reason must be one of those given by DoD Directive 5200.20, Distribution Statements on Technical Documents, as supplemented by the 18 OCT 1983 SECDEF Memo, "Control of Unclassified Technology with Military Application." The Distribution Statement should provide for the broadest distribution possible within limits of security and controlling office limitations.

Block 4. Performing Organization Report Number(s): Enter the unique alphanumeric report number(s) assigned by the organization originating or generating the report from its research and whose name appears in Block 6. These numbers should be in accordance with ANSI STD 239-74, "American National Standard Technical Report Number." If the Performing Organization is also the Monitoring Agency, enter the report number in Block 4.

Block 5. Monitoring Organization Report Number(s): Enter the unique alphanumeric report number(s) assigned by the Monitoring Agency. This should be a number assigned by a DoD or other government agency and should be in accordance with ANSI STD 239-74. If the Monitoring Agency is the same as the Performing Organization, enter the report number in Block 4 and leave Block 5 blank.

Block 6a. Name of Performing Organization: For in-house reports, enter the name of the performing activity. For reports prepared under contract or grant, enter the contractor or the grantee who generated the report and identify the appropriate corporate division, school, laboratory, etc., of the author.

Block 6b. Office Symbol: Enter the office symbol of the Performing Organization.

Block 6c. Address: Enter the address of the Performing Organization. List city, state, and ZIP code.

Block 7a. Name of Monitoring Organization: This is the agency responsible for administering or monitoring a project, contract, or grant. If the monitor is also the Performing Organization, leave Block 7a blank. In the case of joint sponsorship, the Monitoring Organization is determined by advance agreement. It can be either an office, a group, or a committee representing more than one activity, service, or agency.

Block 7b. Address: Enter the address of the Monitoring Organization. Include city, state, and ZIP code.

Block 8a. Name of Funding/Sponsoring Organization: Enter the full official name of the organization under whose immediate funding the document was generated, whether the work was done in-house or by contract. If the Monitoring Organization is the same as the Funding Organization, leave 8a blank.

Block 8b. Office Symbol: Enter the office symbol of the Funding/Sponsoring Organization.

Block 8c. Address: Enter the address of the Funding/Sponsoring Organization. Include city, state and ZIP code.

Block 9. Procurement Instrument Identification Number: For a contractor grantee report, enter the complete contract or grant number(s) under which the work was accomplished. Leave this block blank for in-house reports.

Block 10. Source of Funding (Program Element, Project, Task Area, and Work Unit Number(s): These four data elements relate to the DoD budget structure and provide program and/or administrative identification of the source of support for the work being carried on. Enter the program element, project, task area, work unit accession number, or their equivalents which identify the principal source of funding for the work required. These codes may be obtained from the applicable DoD forms such as the DD Form 1498 (Research and Technology Work Unit Summary) or from the fund citation of the funding instrument. If this information is not available to the authoring activity, these blocks should be filled in by the responsible DoD Official designated in Block 22. If the report is funded from multiple sources, identify only the Program Element and the Project, Task Area, and Work Unit Numbers of the principal contributor.

Block 11. Title: Enter the title in Block 11 in initial capital letters exactly as it appears on the report. Titles on all classified reports, whether classified or unclassified, must be immediately followed by the security classification of the title enclosed in parentheses. A report with a classified title should be provided with an unclassified version if it is possible to do so without changing the meaning or obscuring the contents of the report. Use specific, meaningful words that describe the content of the report so that when the title is machine-indexed, the words will contribute useful retrieval terms.

If the report is in a foreign language and the title is given in both English and a foreign language, list the foreign language title first, followed by the English title enclosed in parentheses. If part of the text is in English, list the English title first followed by the foreign language title enclosed in parentheses. If the title is given in more than one foreign language, use a title that reflects the language of the text. If both the text and titles are in a foreign language, the title should be translated, if possible, unless the title is also the name of a foreign periodical. Transliterations of often used foreign alphabets (see Appendix A of MIL-STD-847B) are available from DTIC in document AD-A080 800.

Block 12. Personal Author(s): Give the complete name(s) of the author(s) in this order: last name, first name, and middle name. In addition, list the affiliation of the authors if it differs from that of the performing organization.

List all authors. If the document is a compilation of papers, it may be more useful to list the authors with the titles of their papers as a contents note in the abstract in Block 19. If appropriate, the names of editors and compilers may be entered in this block.

Block 13a. Type of Report: Indicate whether the report is summary, final, annual, progress, interim, etc.

Block 13b. Time Covered: Enter the inclusive dates (year, month, day) of the period covered, such as the life of a contract in a final contractor report.

Block 14. Date of Report: Enter the year, month, and day, or the year and the month the report was issued as shown on the cover.

Block 15. Page Count: Enter the total number of pages in the report that contain information, including cover, preface, table of contents, distribution lists, partial pages, etc. A chart in the body of the report is counted even if it is unnumbered.

Block 16. Supplementary Notation: Enter useful information about the report in hand, such as: "Prepared in cooperation with..." "Translation of (or by)..." "Symposium..." If there are report numbers for the report which are not noted elsewhere on the form (such as internal series numbers or participating organization report numbers) enter in this block

Block 17. COSATI Codes: This block provides the subject coverage of the report for announcement and distribution purposes. The categories are to be taken from the "COSATI Subject Category List" (DoD Modified), Oct 65, AD-624 000. A copy is available on request to any organization generating reports for DoD. At least one entry is required as follows:

Field - to indicate subject coverage of report.

Group - to indicate greater subject specificity of information in the report.

Sub-Group - if specificity greater than that shown by Group is required, use further designation as the numbers after the period (.) in the Group breakdown. Use only the designation provided by AD-624 000.

Example: The subject "Solid Rocket Motors" is Field 21, Group 08, Subgroup 2 (page 32, AD-624 000).

Block 18. Subject Terms: These may be descriptors, keywords, posting terms, identifiers, open-ended terms, subject headings, acronyms, code words, or any words or phrases that identify the principal subjects covered in the report, and that conform to standard terminology and are exact enough to be used as subject index entries. Certain acronyms or "buzz words" may be used if they are recognized by specialists in the field and have a potential for becoming accepted terms. "Laser" and "Reverse Osmosis" were once such terms.

If possible, this set of terms should be selected so that the terms individually and as a group will remain UNCLASSIFIED without losing meaning. However, priority must be given to specifying proper subject terms rather than making the set of terms appear "UNCLASSIFIED." Each term on classified reports must be immediately followed by its security classification, enclosed in parentheses.

For reference on standard terminology the "DTIC Retrieval and Indexing Terminology" DRIT-1979, AD-A068 500, and the DoD "Thesaurus of Engineering and Scientific Terms (TEST) 1968, AD-672 000, may be useful.

Block 19. Abstract: The abstract should be a pithy, brief (preferably not to exceed 300 words), factual summary of the most significant information contained in the report. However, since the abstract may be machine-searched, all specific and meaningful words and phrases which express the subject content of the report should be included, even if the word limit is exceeded.

If possible, the abstract of a classified report should be unclassified and consist of publicly releasable information (Unlimited), but in no instance should the report content description be sacrificed for the security classification.

NOTE: An unclassified abstract describing a classified document may appear separately from the document in an unclassified context e.g., in DTIC announcement or bibliographic products. This must be considered in the preparation and marking of unclassified abstracts.

For further information on preparing abstracts, employing scientific symbols, verbalizing, etc., see paragraphs 2.1(n) and 2.3(b) in MIL-STD-847B

Block 20. Distribution / Availability of Abstract: This block must be completed for all reports. Check the applicable statement: "unclassified / unlimited," "same as report," or, if the report is available to DTIC registered users "DTIC users."

Block 21. Abstract Security Classification: To ensure proper safeguarding of information, this block must be completed for all reports to designate the classification level of the entire abstract. For CLASSIFIED abstracts, each paragraph must be preceded by its security classification code in parentheses.

Block 22a,b,c. Name, Telephone and Office Symbol of Responsible Individual: Give name, telephone number, and office symbol of DoD person responsible for the accuracy of the completion of this form

AFOSR-TR- 87 - 1983

Air Force Office of Scientific Research
Grant 87-0133

Fast Algorithms for Euler and Navier-Stokes Simulations



Final Technical Report

30 November 1987

*Institute for Scientific Computing
P.O. Box 1988
Fort Collins, Colorado 80522*

Accession For	
NTIS	CRA&I
DTIC	TAB
Unannounced	
Justification	
By	
Date Rec'd	
Availability	
Dist	Author
A-1	

Summary

An explicit flow solver, applicable to the hierarchy of model equations ranging from Euler to full Navier-Stokes, is combined with several techniques designed to reduce computational expense. The computational domain consists of local grid refinements embedded in a global coarse mesh, where the locations of these refinements are defined by the physics of the flow. Flow characteristics are also used to determine which set of model equations is appropriate for solution in each region, thereby reducing not only the number of grid points at which the solution must be obtained, but also the computational effort required to get that solution. Acceleration to steady-state is achieved by applying multigrid on each of the subgrids, regardless of the particular model equations being solved. Since each of these components is explicit, advantage can readily be taken of the vector- and parallel-processing capabilities of machines such as the Cray X-MP and Cray-2.

1. Introduction

It is generally recognized that a comprehensive approach to the simulation of flows involving both complex geometries and complex physics will require powerful advanced-architecture supercomputers with very large memories. Machines capable of producing solutions to Reynolds-averaged Navier-Stokes flows over complex geometries within computing times short enough to be of design interest are expected to be available by the end of this decade [1]. In order to use these parallel-processing supercomputers effectively, algorithms must be adapted to focus the power of multiple processing units on a single flow simulation. Furthermore, the history of computational aerodynamics teaches that the pace of progress in this field is set by the synergism between improved computers and better algorithms. In the past 15 years, improved computers have reduced the cost of computation by a factor of about 100. Over the same period, better algorithms have reduced the cost of computation on a given computer by a factor of almost 1000 [2]. Thus, it is to be expected that the need for faster algorithms will not be diminished by the availability of faster and larger computers.

The most popular algorithms presently in use for calculating three-dimensional Navier-Stokes flows are Beam-Warming [3] (or similar implicit methods) and two types of explicit schemes, Runge-Kutta [4] and Lax-Wendroff [5]. Implicit schemes are highly efficient on uniprocessor machines, and may even be adapted to parallel computers with a small number of processors and shared memory [6]. However, as shown by Bruno [7], they are extremely sensitive to the size and location of memory in large multiprocessor systems. Runge-Kutta and Lax-Wendroff methods, on the other hand, being explicit, map readily onto parallel architectures. The authors chose to use MacCormack's method in the present work because of its robustness and their experience with it, although another explicit scheme, such as Runge-Kutta, could be used in its place. The approach selected enhances the efficiency of the MacCormack scheme by implementing it on a collection of local meshes embedded in a global mesh. Either the Euler, thin-layer Navier-Stokes or full Navier-Stokes equations are solved on designated meshes. The choice of model equations is determined by the nature of the flow physics to be resolved on a particular mesh. When the requirement for time accuracy is relaxed, a convergence acceleration procedure is applied simultaneously to all meshes and all model equations. The entire algorithm is explicit and is designed to perform well on computers consisting of multiple processing units, each having vector processing capability. Examples of such machines are the Cray X-MP and Cray 2.

All of the above-mentioned elements of the algorithm have been integrated into a fully three-dimensional Navier-Stokes flow solver, the performance of which is being evaluated. This task requires a very large memory, high-speed computer such as the Cray-2. It has 256 million words of shared memory and four vector CPUs and is the principal machine being used for development and testing of the scheme. However, until mature, reliable multitasking software is available on the Cray-2, this aspect of the development will be continued on the X-MP, which is also a four-processor system. Since the X-MP presently has at most 16 million words of primary memory, the problem has to be scaled back by decreasing the number of grid points. It is conjectured that parallel-processing efficiency degrades for smaller problems, so that the X-MP will provide a lower bound on expected performance of the full-scale simulation.

2. Equations of Motion

The nondimensional equations of motion may be written in conservation-law form as

$$q_t = -(F_x + G_y + H_z)$$

where, for the Reynolds-averaged Navier-Stokes equations,

$$F = f - Re^{-1}p \quad G = g - Re^{-1}r \quad H = h - Re^{-1}s$$

while, for their thin-layer version,

$$F = f \quad G = g \quad H = h - Re^{-1}d$$

and, for the Euler equations,

$$F = f \quad G = g \quad H = h$$

where:

$$q = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ E \end{bmatrix} \quad f = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uw \\ (E + p)u \end{bmatrix} \quad g = \begin{bmatrix} \rho v \\ \rho uv \\ \rho v^2 + p \\ \rho vw \\ (E + p)v \end{bmatrix}$$

$$h = \begin{bmatrix} \rho w \\ \rho u w \\ \rho v w \\ \rho w^2 + p \\ (E + p)w \end{bmatrix} \quad d = \begin{bmatrix} 0 \\ \mu u_x \\ \mu v_z \\ (\lambda + 2\mu)w_x \\ \gamma \kappa Pr^{-1} e_x + (\lambda + 2\mu)w w_x \end{bmatrix}$$

$$p = \begin{bmatrix} 0 \\ \tau_{zz} \\ \tau_{yz} \\ \tau_{zx} \\ \beta_z \end{bmatrix} \quad r = \begin{bmatrix} 0 \\ \tau_{xy} \\ \tau_{yy} \\ \tau_{zy} \\ \beta_y \end{bmatrix} \quad s = \begin{bmatrix} 0 \\ \tau_{zz} \\ \tau_{yz} \\ \tau_{zx} \\ \beta_z \end{bmatrix}$$

$$\tau_{zz} = \lambda(u_x + v_y + w_z) + 2\mu u_x \quad \beta_z = \gamma \kappa Pr^{-1} e_x + u \tau_{zz} + v \tau_{zy} + w \tau_{zz}$$

$$\tau_{yy} = \lambda(u_x + v_y + w_z) + 2\mu v_y \quad \beta_y = \gamma \kappa Pr^{-1} e_y + u \tau_{yz} + v \tau_{yy} + w \tau_{yz}$$

$$\tau_{zz} = \lambda(u_x + v_y + w_z) + 2\mu w_z \quad \beta_z = \gamma \kappa Pr^{-1} e_z + u \tau_{zz} + v \tau_{zy} + w \tau_{zz}$$

$$\tau_{xy} = \tau_{yx} = \mu(u_y + v_x), \quad \tau_{xx} = \tau_{zz} = \mu(u_x + w_x), \quad \tau_{yz} = \tau_{zy} = \mu(v_z + w_y)$$

Here ρ , u , v , w , p and E are respectively density, velocity components in the x -, y - and z -directions, pressure and total energy per unit volume. This final quantity may be expressed as

$$E = \rho \left(e + \frac{1}{2} (u^2 + v^2 + w^2) \right)$$

where the specific internal energy, e , is related to the pressure and density by the simple law of a calorically-perfect gas

$$p = (\gamma - 1)\rho e$$

with γ denoting the ratio of specific heats. The coefficient of thermal conductivity, κ , and the viscosity coefficients, λ and μ , are assumed to be functions only of temperature. Furthermore, λ is expressed in terms of the dynamic viscosity, μ , by invoking Stokes' assumption of zero bulk viscosity. Re and Pr denote the Reynolds and Prandtl numbers, respectively. Although, for simplicity, the equations of motion are presented here written in Cartesian coordinates, it is well known that their strong conservation law form may be maintained under an arbitrary space- and time-dependent transformation of coordinates.

3. Algorithm Strategy

In order to minimize the cost of simulating complex, three-dimensional viscous flow over complete configurations, the following strategy has been developed:

- a. Use a robust and flexible explicit flow solver capable of simulating either steady or unsteady flow with the Euler, thin-layer Navier-Stokes or full Navier-Stokes equations.
- b. Distribute grid points optimally by making use of both grid stretching and locally-embedded grid refinements.
- c. Make use of a zonal flow simulation strategy ranging from the Euler equations through the full Navier-Stokes equations in order to minimize the computational work per grid point.
- d. Accelerate the convergence of steady flow simulations by means of an explicit multigrid technique which may be applied, without modification, to the entire hierarchy of model equations. Use additional convergence acceleration methods, such as residual averaging, as appropriate.
- e. Take advantage of the explicit nature of the algorithm by mapping it onto a supercomputer architecture consisting of multiple vector-processing CPUs and thus enhance its performance by means of both vectorization and multitasking.

Further detail concerning this strategy is provided in [8].

4. Parallel Processing Considerations

Parallel processing may be viewed in terms of a collection of separately-running programs, called processes, which exchange information among themselves by means of some interconnection scheme. The effects of load balancing, granularity, overhead, and Amdahl's law are all important factors affecting the performance of parallel computers.

Amdahl's law [9] points out that if a computer has two speeds of operation, the slower mode will dominate performance as the faster mode becomes arbitrarily fast. This can be expressed by the relation

$$S = (R + (1 - R) / N)^{-1}$$

where S is the maximum speedup achievable by using N processors on an application which has a fraction of code, R , which must be executed in sequential mode. For example, if R is assumed to be .02, the maximum speedup attainable on a four-processor machine will be 3.77.

The effects on speedup of granularity and overhead, which are also important considerations in multitasking, may be illustrated as follows:

$$S = G (O + G / N)^{-1}$$

Granularity, G , is defined as the length of time required to execute some code segment on a single processor. Although smaller-grained tasks are generally easier to create than larger ones, the overhead, O , associated with creation, synchronization, etc., may negate any performance gains which would otherwise result from parallel execution of smaller tasks. To attain a speedup of 3.77 using four processors would, for example, require that the granularity be more than 65 times the overhead.

A third concern in multitasking is load balancing, or the distribution of computational work across some number of processors. Static and dynamic partitioning may be employed to try to keep all processors equally busy. Static partitioning is most effective when tasks of equal work can be defined *a priori*; dynamic partitioning may enable better load balancing for tasks of varying length, if the additional synchronization overhead incurred is not too great.

Extensions to programming languages which allow the creation and termination of processes, synchronization of processes, and communication among them are necessary for multiprocessing. Both the Cray-2 and the Cray X-MP have software libraries that provide such multitasking tools. Two variations of multitasking are available on the X-MP, namely, macrotasking and microtasking. Macrotasking is intended for application to large-grained problem partitioning, while microtasking, by virtue of its very low overhead, may be used efficiently at a fine-grained level. Only macrotasking is available on the Cray-2. For the current application, both types of multitasking have been examined and tested. Microtasking is easily implemented in a code which has been optimized for vectorization. Macrotasking, however, requires careful examination of the problem in order to define large code structures for parallel execution. More detailed discussions of these multitasking concepts and others may be found in Larson [10], Johnson [11], and Misegades et al. [12].

5. Test Problem

The geometry of the three-dimensional model problem is representative of a turbomachinery application and consists of a rectilinear cascade of finite-span, swept blades mounted between endwalls. Test cases include inviscid subsonic flows and transonic flows with shocks, and viscous laminar and turbulent flows for Reynolds numbers ranging from 8.4×10^3 to 2.0×10^5 (based on cascade gap and critical speed). We believe internal flow problems to be more challenging than external problems for a number of reasons. These include the fact that internal problems limit one's ability to use grid stretching, and that lateral solid boundaries slow convergence by only letting transients propagate out the inlet and exit rather than radiating to infinity in all directions.

The computational domain is partitioned by the collection of embedded meshes. Three levels of grid fineness are used in the present application. If the coarsest mesh (grid 3) is thought of as covering the whole domain, the embedded meshes are then formed by halving or quartering the grid spacing in selected regions. Grid 1 refers to the finest mesh, which lies along the juncture between the blade and the endwalls. Grid 2, coarser than grid 1 by a factor of two, encompasses all surfaces not in grid 1, i.e., the blade and wall surfaces away from the corners. Any coarse-grid points underlying the finer grids are coincident with points on those grids. The intergrid boundaries are treated by overlapping the grids such that the boundary of an embedded mesh lies on the interior of one of its neighboring meshes, with interpolation used where necessary to fill in the surfaces.

When the set of three grids described above is used, the full Navier-Stokes equations are solved on mesh 1, the thin-layer Navier-Stokes equations are solved on mesh 2 and the Euler equations are solved on the coarsest mesh, mesh 3. The flowfield updating begins with mesh 1. After one timestep on mesh 1, mesh 2 is updated exterior to mesh 1 while convergence acceleration is applied at the mesh-2 points interior to mesh 1. Next, mesh 3 is updated exterior to mesh 2 while convergence acceleration is applied at the mesh-3 points interior to mesh 2. Updating proceeds in this fashion until the global mesh has been advanced by one timestep. Then the updating cycle is completed by applying convergence acceleration to coarsenings of the global grid. This cycle is repeated until the desired measure of convergence is satisfied.

6. Performance Evaluation

Performance of parallel computers is evaluated by comparing wall clock time for both uniprocessed and multiprocessed runs on a dedicated machine. This ratio is called the speedup. Dividing the speedup by the number of processors gives the efficiency of processor utilization, a measure of load balancing and overhead.

A two-dimensional version of the code makes extensive use of parallelism inherent in the physical problem to obtain a good load balance when using the X-MP macrotasking software. As the macrotasking approach requires the use of calls to a subroutine library, it has rather high overhead and thus yields best results for large-grained code segments. Table I presents some two-dimensional macrotasking results from both a Cray X-MP/48 and a Cray-2, for the basic solver with multigrid. X-MP performance shows that the algorithm has been efficiently parallelized, while the poorer performance on the Cray-2 is due to less mature macrotasking software on that machine. In Table II, macrotasking results for the basic solver are shown and contrasted with the same code run using the microtasking approach. Microtasking is managed within CPUs, through the use of the X-MP cluster registers. The very low overhead attained by microtasking allows users to partition code at a fine-grained level while still making efficient use of two or more CPUs. The microtasking results in Table II are only marginally better than the macrotasking ones because the algorithm employed in the test had been restructured to maximize task granularity. Three-dimensional microtasking results for a small-grained partitioning of the algorithm are shown in Table III.

Comparison of the embedded-mesh algorithm with a single-mesh algorithm yields the following general conclusion: the accuracy of the embedded-mesh results is essentially that of a global finest mesh, while the convergence rate is like that of a global coarsest mesh. In two-dimensional computations, using the Euler and thin-layer Navier-Stokes equations and three mesh regions, embedding speedups as high as 30 in comparison to a single-mesh algorithm have been obtained (see Table IV). A three-dimensional algorithm has been designed and implemented. As shown in Table V, results have been obtained for simple embeddings which span the y direction, using relatively coarse grids with no tuning. These results are consistent with their two-dimensional analogs.

7. Conclusions

A procedure for solving complex three-dimensional aerodynamic flows on parallel-processing supercomputers has been presented. This procedure incorporates a number of innovations in order to attain high levels of computational efficiency. These innovations include: locally-embedded mesh refinements, a zonal flow simulation strategy that solves the Euler equations through the full Navier-Stokes equations, multigrid convergence acceleration applied to a robust explicit basic flow solver, and both vectorization and multitasking.

Computations have been carried out on parallel-processing supercomputers, principally on the Cray-2, but also on the Cray X-MP because of its more sophisticated multitasking software. The results presented here illustrate that a four CPU shared-memory multiprocessor can be used to carry out aerodynamics simulations with a high degree of efficiency.

The embedded grid scheme has demonstrated performance increases on the order of 30 compared to the global fine grid solution, while maintaining the fine grid solution accuracy.

References

1. Bailey, F.R.: Overview of NASA's Numerical Aerodynamic Simulation Program. Proceedings of the International Symposium on Computational Fluid Dynamics, Sponsored by the Japan Society of Computational Fluid Dynamics, Tokyo, 1985.
2. Peterson, V.L.: Impact of Computers on Aerodynamics Research and Development. Proceedings of the IEEE, Vol. 72, No. 1, January 1984.
3. Pulliam, T.H. and Steger, J.L.: Implicit Finite-Difference Simulation of Three-Dimensional Compressible Flow. AIAA J., Vol. 18, No. 2, pp. 159-167, 1980.
4. Jameson, A., and Baker, T.J.: Solution of the Euler Equations for Complex Configurations. AIAA Paper 83-1929, 1983.

5. MacCormack, R.W.: The Effect of Viscosity in Hypervelocity Impact Cratering. AIAA Paper 69-354, 1969.
6. Chen, S.S.: Advanced Large-Scale and High-Speed Multiprocessor System for Scientific Applications. AIAA Seventh Computational Fluid Dynamics Conference, Cincinnati, Ohio, July, 1985.
7. Bruno, J.: Report on the Feasibility of Hypercube Concurrent Processing Systems in Computational Fluid Dynamics. RIACS TR 86.7, March, 1986.
8. Johnson, G.M, Swisshelm, J.M., Pryor, D.V., and Ziebarth, J.P.: Multitasked Embedded Multigrid for Three-Dimensional Flow Simulation. Lecture Notes in Physics, Vol. 264, Springer, Berlin, 1986, pp. 350-356.
9. Amdahl, G.M.: Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities. Proceedings of the American Federation of Information Processing Societies, Vol. 30, Thompson, Washington, D.C, 1967, pp. 483-485.
10. Larson, J.L.: Practical Concerns in Multitasking on the Cray X-MP. Workshop on Using Multiprocessors in Meteorological Models, ECMWF, Reading, England, 3-6 December 1984.
11. Johnson, G.M.: Parallel Processing in Fluid Dynamics. Proceedings of the ASME Symposium on Parallel-Processing in Fluid Mechanics, FED - Vol. 47, Cincinnati, Ohio, June 1987, pp. 1-8.
12. Misegades, K., Krause, L., and Booth, M.: Microtasking of Fluid Mechanics Codes on the Cray X-MP. Proceedings of the ASME Symposium on Parallel-Processing in Fluid Mechanics, FED - Vol. 47, Cincinnati, Ohio, June 1987, pp. 19-25.

Tables

Table I. Two-Dimensional Macrotasked Multigrid Scheme Performance

Machine	2 Processors		4 Processors	
	Speedup	Efficiency	Speedup	Efficiency
Cray X-MP	1.87	0.94	3.30	0.83
Cray 2	1.80	0.90	2.58	0.65

Table II. Two-Dimensional Multitasked Basic Solver Performance

Machine	2 Processors		4 Processors	
	Speedup	Efficiency	Speedup	Efficiency
Cray X-MP with macrotasking	1.91	0.96	3.58	0.90
Cray X-MP with microtasking	1.93	0.97	3.78	0.95

Table III. Three-Dimensional Microtasked Multigrid Scheme Performance

Machine	2 Processors		3 Processors		4 Processors	
	Speedup	Efficiency	Speedup	Efficiency	Speedup	Efficiency
Cray X-MP	1.96	0.98	2.83	0.94	3.55	0.89

Table IV. Two-Dimensional Embedding Speedups
 (129 x 33 x 33 finest grid, 2 embeddings)

Test Case	Speedup
Inviscid Subcritical	16.4
Inviscid Supercritical	6.1
Turbulent Viscous	30.2

Table V. Three-Dimensional Embedding Speedups
 (65 x 17 x 17 finest grid, 1 embedding)

Test Case	Speedup
Inviscid Subcritical	7.0
Inviscid Supercritical	4.6
Turbulent Viscous	16.0

END

DATE

FILMED

5-88

DTIC